arXiv:1702.05900v1 [cs.CG] 20 Feb 2017

δ-Greedy t-spanner

Gali Bar-On and Paz Carmi

Department of Computer Science,
Ben-Gurion University of the Negev, Israel

Abstract. We introduce a new geometric spanner, δ-*Greedy*, whose construction is based on a generalization of the known *Path-Greedy* and *Gap-Greedy* spanners. The δ-Greedy spanner combines the most desirable properties of geometric spanners both in theory and in practice. More specifically, it has the same theoretical and practical properties as the Path-Greedy spanner: a natural definition, small degree, linear number of edges, low weight, and strong $(1+\varepsilon)$-spanner for every $\varepsilon > 0$. The δ-Greedy algorithm is an improvement over the Path-Greedy algorithm with respect to the number of shortest path queries and hence with respect to its construction time. We show how to construct such a spanner for a set of n points in the plane in $O(n^2 \log n)$ time.
The δ-Greedy spanner has an additional parameter, δ, which indicates how close it is to the Path-Greedy spanner on the account of the number of shortest path queries. For $\delta = t$ the output spanner is identical to the Path-Greedy spanner, while the number of shortest path queries is, in practice, linear.
Finally, we show that for a set of n points placed independently at random in a unit square the expected construction time of the δ-Greedy algorithm is $O(n \log n)$. Our analysis indicates that the δ-Greedy spanner gives the best results among the known spanners of expected $O(n \log n)$ time for random point sets. Moreover, the analysis implies that by setting $\delta = t$, the δ-Greedy algorithm provides a spanner identical to the Path-Greedy spanner in expected $O(n \log n)$ time.

1 Introduction

Given a set P of points in the plane, a Euclidean t-spanner for P is an undirected graph G, where there is a t-spanning path in G between any two points in P. A path between points p and q is a t-spanning path if its length is at most t times the Euclidean distance between p and q (i.e., $t|pq|$).

The most known algorithm for computing t-spanner is probably the *Path-Greedy* spanner. Given a set P of n points in the plane, the Path-Greedy spanner algorithm creates a t-spanner for P as follows. It starts with a graph G having a vertex set P, an empty edge set E and $\binom{n}{2}$ pairs of distinct points sorted in a non-decreasing order of their distances. Then, it adds an edge between p and q to the set E if the length of the shortest path between p and q in G is more than $t|pq|$, see Algorithm 1 for more details. It has been shown in [7, 8, 11, 10, 16, 20] that for every set of points, the Path-Greedy spanner has $O(n)$ edges, a

Algorithm 1 Path-Greedy(P, t)

Input: A set P of points in the plane and a constant $t > 1$
Output: A t-spanner $G(V, E)$ for P

1: sort the $\binom{n}{2}$ pairs of distinct points in non-decreasing order of their distances and store them in list L
2: $E \longleftarrow \emptyset$
3: **for** $(p, q) \in L$ consider pairs in increasing order **do**
4: $\pi \longleftarrow$ length of the shortest path in G between p and q
5: **if** $\pi > t|pq|$ **then**
6: $E := E \cup |pq|$
7: **return** $G = (P, E)$

bounded degree and total weight $O(wt(MST(P)))$, where $wt(MST(P))$ is the weight of a minimum spanning tree of P. The main weakness of the Path-Greedy algorithm is its time complexity – the naive implementation of the Path-Greedy algorithm runs in near-cubic time. By performing $\binom{n}{2}$ shortest path queries, where each query uses Dijkstra's shortest path algorithm, the time complexity of the entire algorithm reaches $O(n^3 \log n)$, where n is the number of points in P. Therefore, researchers in this field have been trying to improve the Path-Greedy algorithm time complexity. For example, the *Approximate-Greedy* algorithm generates a graph with the same theoretical properties as the Path-Greedy spanner in $O(n \log n)$ time [12, 18]. However, in practice there is no correlation between the expected and the unsatisfactory resulting spanner as shown in [13, 15]. Moreover, the algorithm is complicated and difficult to implement.

Another attempt to build a t-spanner more efficiently is introduced in [14, 13]. This algorithm uses a matrix to store the length of the shortest path between every two points. For each pair of points, it first checks the matrix to see if there is a t-spanning path between these points. In case the entry in the matrix for this pair indicates that there is no t-spanning path, it performs a shortest path query and updates the matrix. The authors in [13] have conjectured that the number of performed shortest path queries is linear. This has been shown to be wrong in [4], as the number of shortest path queries may be quadratic. In addition, Bose et al. [4] have shown how to compute the Path-Greedy spanner in $O(n^2 \log n)$ time. The main idea of their algorithm is to compute a partial shortest path and then extend it when needed. However, the drawback of this algorithm is that it is complex and difficult to implement. In [1], Alewijnse et al. compute the Path-Greedy spanner using linear space in $O(n^2 \log^2 n)$ time by utilizing the Path-Greedy properties with respect to the Well Separated Pair Decomposition (WSPD). In [2], Alewijnse et al. compute a t-spanner in $O(n \log^2 n \log^2 \log n)$ expected time by using bucketing for short edges and by using WSPD for long edges. Their algorithm is based on the assumption that the Path-Greedy spanner consists of mostly short edges.

Additional effort has been put in developing algorithms for computing t-spanner graphs, such as θ-Graph algorithm [9, 17], Sink spanner, Skip-List span-

ner [3], and WSPD-based spanners [6, 5]. However, none of these algorithms produces a t-spanner as good as the Path-Greedy spanner in all aspects: size, weight and maximum degree, see [13, 15].

Therefore, our goal is to develop a simple and efficient algorithm that achieves both the theoretical and practical properties of the Path-Greedy spanner. In this paper we introduce the δ-Greedy algorithm that constructs such a spanner for a set of n points in the plane in $O(n^2 \log n)$ time. Moreover, we show that for a set of n points placed independently at random in a unit square the expected running time of the δ-Greedy algorithm is $O(n \log n)$.

2 δ-Greedy

In this section we describe the δ-Greedy algorithm (Section 2.1) for a given set P of points in the plane, and two real numbers t and δ, such that $1 < \delta \leq t$. Then, in Section 2.2 we prove that the resulting graph is indeed a t-spanner with bounded degree. Throughout this section we assume that $\delta < t$ (for example, $\delta = t^{\frac{4}{5}}$ or $\delta = \frac{1+4t}{5}$), except in Lemma 4, where we consider the case that $\delta = t$.

2.1 Algorithm description

For each point $p \in P$ we maintain a collection of cones C_p with the property that for each point $q \in P$ that lies in C_p there is a t-spanning path between p and q in the current graph. The main idea of the δ-Greedy algorithm is to ensure that two cones of a constant angle with apexes at p and q are added to C_p and to C_q, respectively, each time the algorithm runs a shortest path query between points p and q.

The algorithm starts with a graph G having a vertex set P, an empty edge set, and an initially empty collection of cones C_p for each point $p \in P$. The algorithm considers all pairs of distinct points of P in a non-decreasing order of their distances. If $p \in C_q$ or $q \in C_p$, then there is already a t-spanning path that connects p and q in G, and there is no need to check this pair. Otherwise, let d be the length of the shortest path that connects p and q in G divided by $|pq|$. Let $c_p(\theta, q)$ denote the cone with apex at p of angle θ, such that the ray $\overrightarrow{pq}$ is its bisector. The decision whether to add the edge (p, q) to the edge set of G is made according to the value of d. If $d > \delta$, then we add the edge (p, q) to G, a cone $c_p(2\theta, q)$ to C_p, and a cone $c_q(2\theta, p)$ to C_q, where $\theta = \frac{\Pi}{4} - \arcsin(\frac{1}{\sqrt{2} \cdot t})$. If $d \leq \delta$, then we do not add this edge to G, however, we add a cone $c_p(2\theta, q)$ to C_p and a cone $c_q(2\theta, p)$ to C_q, where $\theta = \frac{\Pi}{4} - \arcsin(\frac{d}{\sqrt{2} \cdot t})$.

In Algorithm 2, we give the pseudo-code description of the δ-Greedy algorithm. In Figure 1, we illustrate a cone collection C_p of a point p and how it is modified during the three scenarios of the algorithm. The figure contains the point p, its collection C_p colored in gray, and three points v, u, and w, such that $|pv| < |pu| < |pw|$. Point v lies in C_p representing the first case, where the algorithm does not change the spanner and proceeds to the next pair without

performing a shortest path query. The algorithm runs a shortest path query between p and u, since $u \notin C_p$ (for the purpose of illustration assume $p \notin C_u$). Figure 1(b) describes the second case of the algorithm, where the length of the shortest path between p and u is at most $\delta|pu|$. In this case the algorithm adds a cone to C_p without updating the spanner. Figure 1(c) describes the third case of the algorithm, where the length of the shortest path between p and w is more than $\delta|pw|$. In this case the algorithm adds a cone to C_p and the edge (p, w) to the spanner.

Algorithm 2 δ-Greedy

Input: A set P of points in the plane and two real numbers t and δ s.t. $1 < \delta \leq t$
Output: A t-spanner for P

1: sort the $\binom{n}{2}$ pairs of distinct points in non-decreasing order of their distances (breaking ties arbitrarily) and store them in list L
2: $E \longleftarrow \emptyset$ /* E is the edge set */
3: $C_p \longleftarrow \emptyset \ \ \forall p \in P$ /* C_p is set of cones with apex at p */
4: $G \longleftarrow (P, E)$ /* G is the resulting t-spanner */
5: **for** $(p, q) \in L$ consider pairs in increasing order **do**
6: **if** $(p \notin C_q)$ and $(q \notin C_p)$ **then**
7: $d \longleftarrow$ length of the shortest path in G between p and q divided $|pq|$
8: **if** $d > \delta$ **then**
9: $E \longleftarrow E \cup \{(p, q)\}$
10: $d \longleftarrow 1$
11: $\theta \longleftarrow \frac{\Pi}{4} - \arcsin(\frac{d}{\sqrt{2} \cdot t})$ /* $\frac{1}{\cos\theta - \sin\theta} = \frac{t}{d}$ */
12: $c_p(2\theta, q) \longleftarrow$ cone of angle 2θ with apex at p and bisector $\overrightarrow{pq}$
13: $c_q(2\theta, p) \longleftarrow$ cone of angle 2θ with apex at q and bisector $\overrightarrow{qp}$
14: $C_p \longleftarrow C_p \cup c_p(2\theta, q)$
15: $C_q \longleftarrow C_q \cup c_q(2\theta, p)$
16: **return** $G = (P, E)$

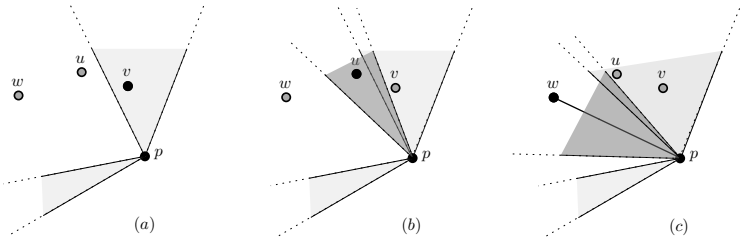

Fig. 1. The three scenarios of the δ-Greedy algorithm. (a) $v \in C_p$; (b) $u \notin C_p$ and $d \leq \delta$; (c) $w \notin C_p$ and $d > \delta$.

2.2 Algorithm analysis

In this section we analyze several properties of the δ-Greedy algorithm, including the spanning ratio and the degree of the resulting graph.

The following lemma is a generalization of Lemma 6.4.1. in [19].

Lemma 1. *Let t and δ be real numbers, such that $1 \leq \delta \leq t$. Let p, q, and r be points in the plane, such that*

1. $p \neq r$,
2. $|pr| \leq |pq|$,
3. $\frac{1}{\cos\theta - \sin\theta} \leq \frac{t}{\delta}$, *where θ is the angle $\angle rpq$* $\left(i.e., \angle rpq = \theta \leq \frac{\Pi}{4} - \arcsin(\frac{\delta}{\sqrt{2}\cdot t})\right)$.

Then $\delta|pr| + t|rq| \leq t|pq|$.

Proof. Let r' be the orthogonal projection of r onto segment $\overline{pq}$. Then, $|rr'| = |pr|\sin\theta$, $|pr'| = |pr|\cos\theta$, and $|r'q| = |pq| - |pr'|$. Thus, $|r'q| = |pq| - |pr|\cos\theta$. By triangle inequality

$$\begin{aligned} |rq| &\leq |rr'| + |r'q| \\ &\leq |pr|\sin\theta + |pq| - |pr|\cos\theta \\ &= |pq| - |pr|(\cos\theta - \sin\theta). \end{aligned}$$

$$\begin{aligned} \text{We have, } \delta|pr| + t|rq| &\leq \delta|pr| + t(|pq| - |pr|(\cos\theta - \sin\theta)) \\ &= t|pq| - t|pr|(\cos\theta - \sin\theta) + \delta|pr| \\ &\leq t|pq| - t|pr|(\cos\theta - \sin\theta) + t(\cos\theta - \sin\theta)|pr| \\ &\leq t|pq|. \end{aligned}$$

□

Lemma 2. *The number of shortest path queries performed by δ-Greedy algorithm for each point is $O(\frac{1}{t/\delta - 1})$.*

Proof. Clearly, the number of shortest path queries performed for each point is at most $n-1$. Thus, we may assume that $t/\delta > 1 + 1/n$. Consider a point $p \in P$ and let (p, q) and (p, r) be two pairs of points that δ-Greedy algorithm has run shortest path queries for. Assume w.l.o.g. that the pair (p, r) has been considered before the pair (p, q), i.e., $|rp| \leq |pq|$. Let d be the length of the path computed by the shortest path query for (p, r) divide by $|pr|$. If $d \leq \delta$, then the cone added to the collection C_p has an angle of at least $\frac{\Pi}{4} - \arcsin(\frac{\delta}{\sqrt{2}\cdot t})$. Otherwise, the algorithm adds the edge (p, r) to G and a new cone to the collection of cones C_p, where the angle of this cone is $\frac{\Pi}{4} - \arcsin(\frac{1}{\sqrt{2}\cdot t})$. Thus, after the shortest path query performed for the pair (p, r), the collection C_p contains a cone $c_p(\theta, r)$, where θ is at least $\frac{\Pi}{2} - 2\arcsin(\frac{\delta}{\sqrt{2}\cdot t})$. The δ-Greedy algorithm performs a shortest path query for (p, q) only if $p \notin C_q$ and $q \notin C_p$. Thus, the angle $\angle rpq$ is at least $\frac{\Pi}{4} - \arcsin(\frac{\delta}{\sqrt{2}\cdot t})$, and we have at most $k = \frac{2\pi}{\theta}$ shortest path queries for a point.

Let us consider the case where $t > 1$ and $\frac{t}{\delta} \to 1$. The equation $\theta = \frac{\Pi}{4} - \arcsin(\frac{\delta}{\sqrt{2}\cdot t})$ implies that $\frac{1}{\cos\theta - \sin\theta} = \frac{t}{\delta}$. Then, we have

$$\theta \to 0,\ \frac{t}{\delta} \sim 1 + \theta,\ \text{and } \theta \sim \frac{t}{\delta} - 1.$$

Thus, we have $k \sim \frac{2\pi}{\frac{t}{\delta}-1} = O(\frac{1}{t/\delta - 1})$. □

Observation 1 *For $\delta = t^{\frac{x-1}{x}}$, where $x > 1$ is a fixed integer, the number of shortest path queries performed by δ-Greedy algorithm for each point is $O(\frac{x}{t-1})$.*

Proof. As in Lemma 2, let us consider the case where $t > 1$ and $\frac{t}{\delta} \to 1$. Then, we have

$$\theta \to 0,\ \ \frac{t}{\delta} \sim 1 + \theta,\ \ \frac{t}{t^{(\frac{x-1}{x})}} \sim 1 + \theta,\ \ t^{(\frac{1}{x})} \sim 1 + \theta,$$

$$t \sim (1+\theta)^x,\ \ t \sim 1 + x \cdot \theta,\ \text{and } \theta \sim \frac{t-1}{x}.$$

Thus, we have $k \sim \frac{2\pi x}{t-1} = O(\frac{x}{t-1})$.

□

Observation 2 *The running time of δ-Greedy algorithm is $O(\frac{n^2 \log n}{(t/\delta - 1)^2})$.*

Proof. First, the algorithm sorts the $\binom{n}{2}$ pairs of distinct points in non-decreasing order of their distances, this takes $O(n^2 \log n)$ time. A shortest path query is done by Dijkstra's shortest path algorithm on a graph with $O(\frac{n}{t/\delta - 1})$ edges and takes $O(\frac{n}{t/\delta - 1} + n \log n)$ time. By Lemma 2 each point performs $O(\frac{1}{t/\delta - 1})$ shortest path queries. Therefore, we have that the running time of δ-Greedy algorithm is $O((\frac{n}{t/\delta - 1})^2 \log n)$. □

Observation 3 *The number of cones that each point has in its collection along the algorithm is constant depending on t and δ ($O(\frac{1}{t/\delta - 1})$).*

Proof. As shown in Lemma 2, the number of shortest path queries for each point is $O(\frac{1}{t/\delta - 1})$. The subsequent step of a shortest path query is the addition of two cones, meaning that for each point p the number of cones in the collection of cones C_p is $O(\frac{1}{t/\delta - 1})$. □

Corollary 1. *The additional space for each point p for the collection C_p is constant.*

Lemma 3. *The output graph $G = (P, E)$ of δ-Greedy algorithm (Algorithm 2) is a t-spanner for P (for $1 < \delta < t$).*

Proof. Let $G = (P, E)$ be the output graph of the δ-Greedy algorithm. To prove that G is a t-spanner for P we show that for every pair $(p, q) \in P$, there exists a t-spanning path between them in G. We prove the above statement by induction

on the rank of the distance $|pq|$, i.e., the place of (p, q) in a non-decreasing distances order of all pairs of points in P.
Base case: Let (p, q) be the first pair in the ordered list (i.e., the closest pair). The edge (p, q) is added to E during the first iteration of the loop in step 9 of Algorithm 2, and thus there is a t-spanning path between p and q in G.
Induction hypothesis: For every pair (r, s) that appears before the pair (p, q) in the ordered list, there is a t-spanning path between r and s in G.
The inductive step: Consider the pair (p, q). We prove that there is a t-spanning path between p and q in G. If $p \notin C_q$ and $q \notin C_p$, we check whether there is a δ-spanning path in G between p and q. If there is a path which length is at most $\delta|pq|$, then $\delta|pq| \leq t|pq|$, meaning there is a t-spanning path between p and q in G. If there is no path of length of at most $\delta|pq|$, we add the edge (p, q) to G, which forms a t-spanning path.

Consider that $p \in C_q$ or $q \in C_p$, and assume w.l.o.g. that $q \in C_p$. Let (p, r) be the edge handled in Step 5 in Algorithm 2 when the cone containing q has been added to C_p (Step 12 in Algorithm 2). Notice that $|pr| \leq |pq|$. Step 7 of Algorithm 2 has computed the value d for the pair (p, r). In the algorithm there are two scenarios depending on the value of d.

The first scenario is when $d > \delta$, then the algorithm has added the edge (p, r) to G and a cone $c_p(\theta, r)$ to C_p, where $\theta = 2(\frac{\Pi}{4} - \arcsin(\frac{1}{\sqrt{2}\cdot t}))$. Thus, the angle between (p, q) and (p, r) is less than $\theta/2$. Hence, $|rq| < |pq|$ and by the induction hypothesis there is a t-spanning path between r and q. Consider the shortest path between p and q that goes through the edge (p, r). The length of this path is at most $|pr| + t|rq|$. By Lemma 1, we have $|pr| + t|rq| \leq \delta|pr| + t|rq| \leq t|pq|$ for $\delta = 1$. Therefore, we have a t-spanning path between p and q.

The second scenario is when $d \leq \delta$, then the algorithm has added a cone $c_p(\theta, r)$ to C_p, where $\theta = 2(\frac{\Pi}{4} - \arcsin(\frac{d}{\sqrt{2}\cdot t}))$. Thus, the angle between (p, q) and (p, r) is less than $\theta/2$. Hence, $|rq| < |pq|$ and by the induction hypothesis there is a t-spanning path between r and q. Consider the shortest path between p and q that goes through r. The length of this path is at most $d|pr| + t|rq|$. By Lemma 1, we have $d|pr| + t|rq| \leq t|pq|$. Therefore, we have a t-spanning path between p and q. □

Theorem 4. *The δ-Greedy algorithm computes a t-spanner for a set of points P with the same properties as the Path-Greedy t-spanner, such as degree and weight, in $O((\frac{n}{t/\delta-1})^2 \log n)$ time.*

Proof. Clearly, the degree of the δ-Greedy is at most the degree of the Path-Greedy δ-spanner. The edges of the δ-Greedy spanner satisfy the δ-leap frog property, thus, the weight of the δ-Greedy is as Path-Greedy t-spanner. Hence, we can pick δ close to t, such that we will have the required bounds. □

Lemma 4. *If $t = \delta$, the result of the δ-Greedy algorithm is identical to the result of the Path-Greedy algorithm.*

Proof. Assume towards contradiction that for $t = \delta$ the resulting graph of the δ-Greedy algorithm, denoted as $G = (P, E)$, differs from the result of the Path-Greedy algorithm, denoted as $G' = (P, E')$. Assuming the same order of the

sorted edges, let (p, q) be the first edge that is different in G and G'. Notice that δ-Greedy algorithm decides to add the edge (p, q) to G when there is no t-spanning path between p and q in G. Since until handling the edge (p, q) the graphs G and G' are identical, the Path-Greedy algorithm also decides to add the edge (p, q) to G'. Therefore, the only case we need to consider is $(p, q) \in E'$ and $(p, q) \notin E$. The δ-Greedy algorithm does not add an edge (p, q) to G in two scenarios:

- there is a t-spanning path between p and q in the current graph G – which contradicts that the Path-Greedy algorithm adds the edge (p, q) to G';
- $p \in C_q$ or $q \in C_p$ – the δ-Greedy algorithm does not perform a shortest path query between p and q. Assume w.l.o.g., $q \in C_p$, and let (p, r) be the edge considered in Step 5 in Algorithm 2 when the cone containing q has been added to C_p. The angle of the added cone is $\theta = \frac{\Pi}{2} - 2\arcsin(\frac{d}{\sqrt{2}\cdot t})$, where d is the length of the shortest path between p and r divided $|pr|$. Thus, we have $|pr| \leq |pq|$ and $\frac{1}{\cos\alpha - \sin\alpha} \leq \frac{t}{d}$, where $\alpha \leq \theta$ is the angle $\angle rpq$. Then, by Lemma 1, $\delta|pr| + t|rq| \leq t|pq|$, and since there is a path from p to r of length at most $\delta|pr|$, we have that there is t-spanning path between p and q in the current graph. This is in contradiction to the assumption that the Path-Greedy algorithm adds the edge (p, q) to E'.

□

3 δ-Greedy in Expected $O(n \log n)$ Time for Random Set

In this section we show how a small modification in the implementation improves the running time of the δ-Greedy algorithm. This improvement yields an expected $O(n \log n)$ time for random point sets. The first modification is to run the shortest path query between points p to q up to $\delta|pq|$. That is, running Dijkstras shortest path algorithm with source p and terminating as soon as the minimum key in the priority queue is larger than $\delta|pq|$.

Let P be a set of n points in the plane uniformly distributed in a unit square. To prove that δ-Greedy algorithm computes a spanner for P in expected $O(n \log n)$ time, we need to show that:

- each point runs a constant number of shortest path queries – follows from Lemma 2;
- the expected number of points visited in each query is constant – The fact that the points are randomly chosen uniformly in the unit square implies that the expected number of points at distance of at most r from point p is $\Theta(r^2 \cdot n)$. A shortest path query from a point p to a point q terminates as soon as the minimum key in the priority queue exceeds $\delta|pq|$, thus, it is expected to visit $O(n \cdot (\delta|pq|)^2)$ points.
  By Lemma 2 the number of shortest path queries performed by the algorithm for a point p is $O(\frac{1}{t/\delta - 1})$. Each such query defines a cone with apex at p of angle $\Omega(t/\delta - 1)$, such that no other shortest path query from p will be

performed to a point in this cone. By picking $k = \frac{1}{t/\delta-1}$ and $r = \frac{k}{\sqrt{n}}$, we have that the expected number of points around each point in a distance of r is $\Theta(k^2) = \Theta(\frac{1}{(t/\delta-1)^2})$.

Assume we partition the plane into k equal angle cones with apex at point p. The probability that there exists a cone that does not contain a point from the set of points of distance $\frac{k}{\sqrt{n}}$ is at most $k \cdot (1 - \frac{1}{k})^{k^2}$. Let Q be the set of points that p computed a shortest path query to, and let $q \in Q$ be the farthest point in Q from p. Then, the expected Euclidean distance between p and q is less than $\frac{k}{\sqrt{n}}$. Thus, the expected number of points visited by the entire set of shortest path queries from a point is $O(\frac{\delta^2 k^2}{t/\delta-1}) = O(\frac{\delta^2}{(t-\delta)^3})$;

- the next pair to be processed can be obtained in expected $O(\log n)$ time without sorting all pairs of distinct points – Even-though this is quite straight forward, for completeness we give a short description how this can be done. Divide the unit square to $n \times n$ grid cells of side length $1/n$. A hash table of size $3n$ is initialized, and for each non-empty grid cell (at most n such cells) we map the points in it to the hash table. In addition, we maintain a minimum heap H_p for each point $p \in P$ (initially empty), and one main minimum heap H that contains the top element of each H_p. Each heap H_p contains a subset of the pairs that include p.

  For each point $p \in P$, all the cells of distance at most $\frac{k}{\sqrt{n}}$ from p are scanned (using the hash table) to find all the points in these cells, where k is a parameter that we fix later. All the points found in these cells are added to H_p according to their Euclidean distance from p.

  The heap H holds the relevant pairs in an increasing order, therefore the pairs are extracted from the main heap H. After extracting the minimum pair in H that belongs to a point p, we add to H the next minimum in H_p. To insure the correctness of the heaps, when needed we increase the distance to the scanned cells. Observe that there may be a pair (p, q) such that $|pq| < |rw|$, where the pair (r, w) is the top pair in H. This can occur only when the pair (p, q) has not been added to H_p nor H_q, and this happens when $p \in C_q$ or $q \in C_p$. However, in this case we do not need to consider the pair (p, q).

  Notice that the only cells that are not contained in C_p are scanned to add more pairs to H_p. Thus, points that are in C_p are ignored.

Therefore, the total expected running time of the algorithm is $O(\frac{\delta^2}{(t-\delta)^3} n \log n)$. Since both t and t/δ are constants bigger than one, the expected running time of the δ-Greedy algorithm is $O(n \log n)$.

A very nice outcome of δ-Greedy algorithm and its analysis can be seen when δ is equal to t. Assume that δ-Greedy algorithm (for $\delta = t$) has computed a shortest path query for two points p and q and the length of the received path is $d|pq|$. If the probability that $t/d > 1 + \varepsilon$ is low (e.g, less than $1/2$), for some constant $\varepsilon > 0$, then δ-Greedy algorithm computes the Path-Greedy spanner with linear number of shortest path queries. Thus δ-Greedy algorithm computes

the Path-Greedy spanner for a point set uniformly distributed in a square in expected $O(n \log n)$ time.

Not surprisingly our experiments have shown that this probability is indeed low (less than 1/100), since most of the shortest path queries are performed on pairs of points placed close to each other (with respect to Euclidean distance), and thus with a high probability their shortest path contains a constant number of points. Moreover, it seems that for a "real-life" input this probably is low. Thus, there is a very simple algorithm to compute the Path-Greedy spanner in expected $O(n^2 \log n)$ time for real-life inputs, based on the δ-Greedy algorithm

For real-life input we mean that our analysis suggests that in the current computers precision (Memory) one cannot create an instance of points set with more than 1000 points, where the Path-Greedy spanner based on the δ-Greedy algorithm has more than $O(n^2 \log n)$ constructing time.

4 Experimental Results

In this section we discuss the experimental results by considering the properties of the graphs generated by different algorithms and the number of shortest path queries performed during these algorithms. We have implemented the Path-Greedy, δ-Greedy, Gap-Greedy, θ-Graph, Path-Greedy on θ-Graph algorithms. The Path-Greedy on θ-graph t-spanner, first computes a θ-graph t'-spanner, where $t' < t$, and then runs the Path-Greedy t/t'-spanner on this t'-spanner. The shortest path queries criteria is used for an absolute running time comparison that is independent of the actual implementation. The known theoretical bounds for the algorithms can be found in Table 1.

The experiments were performed on a set of 8000 points, with different values of the parameter δ (between 1 and t). We have chosen to present the parameter δ for the values $t, t^{0.9}$ and $\sqrt{t}$. This values do not have special properties, they where chosen arbitrary to present the behavior of the spanner.

To avoid the effect of specific instances, we have run the algorithms several times and taken the average of the results. However, in all the cases the difference between the values is negligible. Table 2–4 show the results of our experiments for different values of t and δ. The columns of the weight (divided by $wt(MST)$) and the degree are rounded to integers, and the columns of the edges are rounded to one digit after the decimal point (in k).

Algorithm	**Edges**	**Weight** $\frac{}{wt(MST)}$	**Degree**	**Time**
Path-Greedy	$O(\frac{n}{t-1})$	$O(1)$	$O(\frac{1}{t-1})$	$O(n^3 \log n)$
Gap-Greedy	$O(\frac{n}{t-1})$	$O(\log n)$	$O(\frac{1}{t-1})$	$O(n \log^2 n)$
θ-Graph	$O(\frac{n}{\theta})$	$O(n)$	$O(n)$	$O(\frac{n}{\theta} \log n)$
δ-Greedy	$O(\frac{n}{t/\delta-1})$	$O(1)$	$O(\frac{1}{t/\delta-1})$	$O(\frac{1}{t/\delta-1} \cdot n^2 \log n)$

Table 1. Theoretical bounds of different t-spanner algorithms

Algorithm	δ	**Edges** (in K)	**Weight** $\frac{}{wt(MST)}$	**Degree**	**Shortest path queries** (in K)
Path-Greedy	-	35.6	10	17	31996
δ-Greedy	1.1	35.6	10	17	254
δ-Greedy	1.0896	37.8	12	18	242
δ-Greedy	1.048	51.6	19	23	204
θ-Graph	-	376.6	454	149	-
Greedy on θ-Graph	1.0896	37.8	12	18	3005
Greedy on θ-Graph	1.048	52	19	23	693
Gap-Greedy	-	51.6	19	23	326

Table 2. Comparison between several t-spanner algorithms for $t = 1.1$

Algorithm	δ	**Edges** (in K)	**Weight** $\frac{}{wt(MST)}$	**Degree**	**Shortest path queries** (in K)
Path-Greedy	-	15.1	3	7	31996
δ-Greedy	1.5	15.1	3	7	82
δ-Greedy	1.44	16	3	8	77
δ-Greedy	1.224	22.5	5	11	63
θ-Graph	-	118.6	76	53	-
Greedy on θ-Graph	1.44	16	3	8	817
Greedy on θ-Graph	1.224	22.5	6	11	198
Gap-Greedy	-	22.6	5	11	95

Table 3. Comparison between several t-spanner algorithms for $t = 1.5$

Algorithm	δ	**Edges** (in K)	**Weight** $\frac{}{wt(MST)}$	**Degree**	**Shortest path queries** (in K)
Path-Greedy	-	11.4	2	5	31996
δ-Greedy	2	11.4	2	5	55
δ-Greedy	1.866	11.9	2	5	52
δ-Greedy	1.414	16.3	3	8	44
θ-Graph	-	85.3	48	42	-
Greedy on θ-Graph	1.866	11.9	3	6	493
Greedy on θ-Graph	1.414	16.5	3	8	129
Gap-Greedy	-	16	3	8	63

Table 4. Comparison between several t-spanner algorithms for $t = 2$

4.1 Implementation details

All the algorithms mentioned above were implemented in Java using JGraphT and JGraph libraries. The experiments were performed on an Intel Xeon CPU E5-2680 v2 @ 2.80 GHz (2 processors) and 128 GB RAM on Windows Server 2012 Standard OS using ECJ for compilation. The sample point sets were generated by Java.util.Random pseudo random number generator.

4.2 Results analysis

The experiments indicate that the δ-Greedy algorithm achieves good results in practice as expected. The outcome of the δ-Greedy algorithm for all values of δ, that have been checked, is roughly the same as the results of the Path-Greedy algorithm for all parameters. Compared to other algorithms, the δ-Greedy graphs are superior to the graphs produced by the n^2-Gap algorithm, and are as good as Path-Greedy on θ-Graph, with significantly a lower number of shortest path queries. The theoretical complexity of the Path-Greedy on θ-Graph is $O(n^2 \log n)$, same as the δ-Greedy algorithm. However in practice the δ-Greedy algorithm computes considerably less shortest path queries. Hence, the δ-Greedy algorithm has the same results in weight, size and degree as the Path-Greedy on θ-Graph algorithm with better running time.

In addition, Farshi and Gudmundsson in [15] have implemented various spanner algorithms and shown that the Path-Greedy algorithm for $t = 1.1$ and for $t = 2$ on random graphs are almost identical to ours experimental results in weight, size and degree. Moreover, they have shown that Path-Greedy spanner is the highest quality geometric spanner in terms of edge count, degree and weight. They have presented the results for $t = 1.1$ and for $t = 2$ on random point set with 8000 points. Moreover, they have shown that the θ-Graph spanner achieves in practice the best results after the Path-Greedy spanner for all parameters that have been tested (size, weight and degree) comparing to other spanners that they have implemented (such as the Approximate-Greedy, the WSPD-spanner, Skip-list and Sink-Spanner). Our experiments show that the δ-Greedy spanner achieves better results than the θ-Graph spanner. Thus, combining this with the results in [15], we conclude that the δ-spanner achieves the highest quality geometric spanner with respect to θ-Graph, Approximate-Greedy, the WSPD-spanner, Skip-list, Sink-Spanner, and Gap-Greedy spanners.

The experiments reinforce the analysis that picking δ very close to t (for example $\delta = t^{0.9}$), the results are very close to the Path-Greedy spanner, and the number of the performed shortest paths queries is still small. Moreover, the experiments show that the number of shortest path queries is linear while selecting $\delta = t$ and obtaining the δ-Greedy spanner identical to the Path-Greedy t-spanner.

The experiments presented in this paper were performed on set of points placed independently at random in a unit square. However, we conjecture that the δ-Greedy algorithm computes a t-spanner in expected $O(n \log n)$ time for almost all realistic inputs, that is, the δ-Greedy algorithm computes a t-spanner

in expected $O(n \log n)$ time for point sets that are not deliberately hand-made to cause a higher number of shortest path queries.

5 Acknowledgments

We would like to thank Rachel Saban for implementing the algorithms.

References

1. S. P. A. Alewijnse, Q. W. Bouts, A. P. ten Brink, and K. Buchin. Computing the greedy spanner in linear space. *Algorithmica*, 73(3):589–606, 2015.
2. S. P. A. Alewijnse, Q. W. Bouts, Alex P. ten Brink, and K. Buchin. Distribution-sensitive construction of the greedy spanner. *Algorithmica*, pages 1–23, 2016.
3. S. Arya, D. M. Mount, and M. H. M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *FOCS*, pages 703–712, 1994.
4. Prosenjit Bose, Paz Carmi, Mohammad Farshi, Anil Maheshwari, and Michiel H. M. Smid. Computing the greedy spanner in near-quadratic time. In *SWAT*, pages 390–401, 2008.
5. P. B. Callahan and S. R. Kosaraju. A decomposition of multi-dimensional point-sets with applications to k-nearest-neighbors and n-body potential fields. In *STOC*, pages 546–556, 1992.
6. Paul B. Callahan. Optimal parallel all-nearest-neighbors using the well-separated pair decomposition. In *FOCS*, pages 332–340, 1993.
7. B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. *Int. J. Comp. Geom. and Applic.*, 5:125–144, 1995.
8. Barun Chandra. Constructing sparse spanners for most graphs in higher dimensions. *Inf. Process. Lett.*, 51(6):289–294, 1994.
9. Kenneth L. Clarkson. Approximation algorithms for shortest path motion planning. In *STOC*, pages 56–65, 1987.
10. G. Das, P. J. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *SoCG*, pages 53–62, 1993.
11. G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. J. Comp. Geom. and Applic.*, 7(4):297–315, 1997.
12. Gautam Das and Giri Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. J. Comp. Geom. and Applic.*, 7(4):297–315, 1997.
13. M. Farshi and J. Gudmundsson. Experimental study of geometric t-spanners: A running time comparison. In *WEA*, pages 270–284, 2007.
14. Mohammad Farshi and Joachim Gudmundsson. Experimental study of geometric t-spanners. In *ESA*, pages 556–567, 2005.
15. Mohammad Farshi and Joachim Gudmundsson. Experimental study of geometric t-spanners. *ACM Journal of Experimental Algorithmics*, 14, 2009.
16. Joachim Gudmundsson, Christos Levcopoulos, and Giri Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31(5):1479–1500, 2002.
17. J. Mark Keil. Approximating the complete Euclidean graph. In *SWAT*, pages 208–213, 1988.
18. Christos Levcopoulos, Giri Narasimhan, and Michiel H. M. Smid. Improved algorithms for constructing fault-tolerant spanners. *Algorithmica*, 32(1):144–156, 2002.

19. Giri Narasimhan and Michiel Smid. *Geometric Spanner Networks*. Cambridge University Press, New York, NY, USA, 2007.
20. José Soares. Approximating Euclidean distances by small degree graphs. *Discrete & Computational Geometry*, 11(2):213–233, 1994.